

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login
 Search: © The ACM Digital Library © The Guide
 "execution plans" and "ranking operators" and "global rank"

THE ACM DIGITAL LIBRARY  Feedback Report a problem Satisfaction survey

Terms used [execution plans](#) and [ranking operators](#)

Found 336 of 158,639

Sort results by

 relevance

 Save results to a Binder

[Try an Advanced Search](#)

Display results

 expanded form

 Search Tips
 Open results in a new window

[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale 

1 Research sessions: non-standard query processing: Rank-aware query optimization 

Ihab F. Ilyas, Rahul Shah, Walid G. Aref, Jeffrey Scott Vitter, Ahmed K. Elmagarmid

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(281.93 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Ranking is an important property that needs to be fully supported by current relational query engines. Recently, several rank-join query operators have been proposed based on rank aggregation algorithms. Rank-join operators progressively rank the join results while performing the join operation. The new operators have a direct impact on traditional query processing and optimization. We introduce a rank-aware query optimization framework that fully integrates rank-join operators into relational qu ...

2 Research papers: optimization: RankSQL: query algebra and optimization for relational top-k queries 

Chengkai Li, Kevin Chen-Chuan Chang, Ihab F. Ilyas, Sumin Song

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(741.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper introduces RankSQL, a system that provides a systematic and principled framework to support efficient evaluations of ranking (*top-k*) queries in relational database systems (RDBMS), by extending relational algebra and query optimization. Previously, *top-k* query processing is studied in the middleware scenario or in RDBMS in a "piecemeal" fashion, *i.e.*, focusing on specific operator or sitting outside the core of query engines. In contrast, we aim to support ranking ...

3 Efficient mid-query re-optimization of sub-optimal query execution plans 

Navin Kabra, David J. DeWitt

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2

Full text available:  pdf(1.83 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

For a number of reasons, even the best query optimizers can very often produce sub-optimal query execution plans, leading to a significant degradation of performance. This is especially true in databases used for complex decision support queries and/or object-relational databases. In this paper, we describe an algorithm that detects sub-optimality of a query execution plan during query execution and attempts to correct the problem. The

basic idea is to collect statistics at key points durin ...

4 Counting, enumerating, and sampling of execution plans in a cost-based query optimizer

Florian Waas, César Galindo-Legaria

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2

Full text available:  pdf(471.32 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Testing an SQL database system by running large sets of deterministic or stochastic SQL statements is common practice in commercial database development. However, code defects often remain undetected as the query optimizer's choice of an execution plan is not only depending on the query but strongly influenced by a large number of parameters describing the database and the hardware environment. Modifying these parameters in order to steer the optimizer to select other plans is difficult since ...

5 Query optimization in the presence of limited access patterns

Daniela Florescu, Alon Levy, Ioana Manolescu, Dan Suciu

June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data**, Volume 28 Issue 2

Full text available:  pdf(1.66 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We consider the problem of query optimization in the presence of limitations on access patterns to the data (i.e., when one must provide values for one of the attributes of a relation in order to obtain tuples). We show that in the presence of limited access patterns we must search a space of annotated query plans, where the annotations describe the inputs that must be given to the plan. We describe a theoretical and experimental analysis of the resulting search space and a ...

6 Web engineering: Quality driven web services composition

Liangzhao Zeng, Boualem Benatallah, Marlon Dumas, Jayant Kalagnanam, Quan Z. Sheng
May 2003 **Proceedings of the 12th international conference on World Wide Web**

Full text available:  pdf(245.01 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The process-driven composition of Web services is emerging as a promising approach to integrate business applications within and across organizational boundaries. In this approach, individual Web services are federated into composite Web services whose business logic is expressed as a process model. The tasks of this process model are essentially invocations to functionalities offered by the underlying component services. Usually, several component services are able to execute a given task, alth ...

Keywords: QoS, service composition, web services

7 Research sessions: query progress: Estimating progress of execution for SQL queries

Surajit Chaudhuri, Vivek Narasayya, Ravishankar Ramamurthy

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(201.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

Today's database systems provide little feedback to the user/DBA on how much of a SQL query's execution has been completed. For long running queries, such feedback can be very useful, for example, to help decide whether the query should be terminated or allowed to run to completion. Although the above requirement is easy to express, developing a robust

indicator of progress for query execution is challenging. In this paper, we study the above problem and present techniques that can form the basi ...

8 Query processing and optimization: Query processing in a geographic mediation system

Mehdi Essid, Omar Boucelma, François-Marie Colonna, Yassine Lassoued

November 2004 **Proceedings of the 12th annual ACM international workshop on Geographic information systems**

Full text available:  pdf(161.67 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Despite various interoperability recommendations, heterogeneity of Geographic Information Systems (GIS) is still an issue. This led to an increasing need for a data integration system that allows transparent and uniform access to spatial data disseminated over a network. In this paper, we describe the internals of query processing in the VirGIS mediation system. Recall that a data mediation system provides users with a uniform access to a multitude of (local/remote) data sources, without dupl ...

Keywords: GML, WFS, data integration, query rewriting

9 A graph-theoretic model for optimizing queries involving methods

Chiang Lee, Chi-Sheng Shih, Yaw-Huei Chen

April 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 4

Full text available:  pdf(266.86 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Traditional algorithms for optimizing the execution order of joins are no more valid when selections and projections involve methods and become very expensive operations. Selections and projections could be even more costly than joins such that they are pulled above joins, rather than pushed down in a query tree. In this paper, we take a fundamental look at how to approach query optimization from a top-down design perspective, rather than trying to force one model to fit into another. We present ...

Keywords: Graph model, Method query, Object-oriented databases, Query optimization, Spanning tree

10 Optimizing SQL queries for parallel execution

G. von Bultzingsloewen

December 1989 **ACM SIGMOD Record**, Volume 18 Issue 4

Full text available:  pdf(448.62 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The optimization problem discussed in this paper is the translation of an SQL query into an efficient parallel execution plan for a multiprocessor database machine under the performance goal of reduced response times as well as increased throughput in a multiuser environment. We describe and justify the most important research problems which have to be solved to achieve this task, and we explain our approach to solve these problems.

11 Efficient and accurate cost models for parallel query optimization (extended abstract)

Sumit Ganguly, Akshay Goel, Avi Silberschatz

June 1996 **Proceedings of the fifteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Full text available:  pdf(1.06 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

An adaptive software library for fast Fourier transforms

Dragan Mirković, Rishad Mahasoom, Lennart Johnsson

May 2000 Proceedings of the 14th international conference on SupercomputingFull text available:  [pdf\(1.15 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we present an adaptive and portable software library for the fast Fourier transform (FFT). The library consists of a number of composable blocks of code called codelets, each computing a part of the transform. The actual FFT algorithm used by the code is determined at run-time by selecting the fastest strategy among all possible strategies, given available codelets, for a given transform size. We also present an efficient automatic method of generating the lib ...

13 Checking the temporal integrity of interactive multimedia documents

I. Mirbel, B. Pernici, T. Sellis, S. Tserkezoglou, M. Vaziriannis

July 2000 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 9 Issue 2Full text available:  [pdf\(269.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

When authoring multimedia scenarios, and in particular scenarios with user interaction, where the sequence and time of occurrence of interactions is not predefined, it is difficult to guarantee the consistency of the resulting scenarios. As a consequence, the execution of the scenario may result in unexpected behavior or inconsistent use of media. The present paper proposes a methodology for checking the temporal integrity of interactive multimedia document (IMD) scenarios at authoring ti ...

Keywords: Constraint networks, Multimedia presentation, Temporal integrity**14 Using local optimality criteria for efficient information retrieval with redundant information filters**

Neil C. Rowe

April 1996 ACM Transactions on Information Systems (TOIS), Volume 14 Issue 2Full text available:  [pdf\(2.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We consider information retrieval when the data—for instance, multimedia—is computationally expensive to fetch. Our approach uses “information filters” to considerably narrow the universe of possibilities before retrieval. We are especially interested in redundant information filters that save time over more general but more costly filters. Efficient retrieval requires that decisions must be made about the necessity, order, and concurrent processing of proposed filte ...

Keywords: Boolean algebra, conjunction, filters, natural language, optimization, queries**15 Research sessions: query processing I: Exploiting statistics on query expressions for optimization**

Nicolas Bruno, Surajit Chaudhuri

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of dataFull text available:  [pdf\(1.33 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Statistics play an important role in influencing the plans produced by a query optimizer. Traditionally, optimizers use statistics built over base tables and assume independence between attributes while propagating statistical information through the query plan. This approach can introduce large estimation errors, which may result in the optimizer choosing

inefficient execution plans. In this paper, we show how to extend a generic optimizer so that it also exploits statistics built on expression ...

16 Supporting top-k join queries in relational databases

F. Ilyas, G. Aref, K. Elmagarmid

September 2004 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 13 Issue 3

Full text available:  pdf(317.70 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Ranking queries, also known as top- k queries, produce results that are ordered on some computed score. Typically, these queries involve joins, where users are usually interested only in the top- k join results. Top- k queries are dominant in many emerging applications, e.g., multimedia retrieval by content, Web databases, data mining, middlewares, and most information retrieval applications. Current relational query processors do not handle ranking queries efficiently, especia ...

Keywords: Query operators, Rank aggregarion, Ranking, Top- k queries

17 Answering queries using views: A survey

Alon Y. Halevy

December 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 4

Full text available:  pdf(308.74 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The problem of answering queries using views is to find efficient methods of answering a query using a set of previously defined materialized views over the database, rather than accessing the database relations. The problem has recently received significant attention because of its relevance to a wide variety of data management problems. In query optimization, finding a rewriting of a query using a set of materialized views can yield a more efficient query execution plan. To support the separat ...

Keywords: Data integration, Date warehousing, Materialized views, Query optimization, Survey, Web-site management

18 The knowledge grid

Mario Cannataro, Domenico Talia

January 2003 **Communications of the ACM**, Volume 46 Issue 1

Full text available:  pdf(109.40 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
 html(26.74 KB)

Designing, building, and implementing an architecture for distributed knowledge discovery.

19 Research sessions: continuous queries and streams: Rate-based query optimization for streaming information sources

Stratis D. Viglas, Jeffrey F. Naughton

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(1.11 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Relational query optimizers have traditionally relied upon table cardinalities when estimating the cost of the query plans they consider. While this approach has been and continues to be successful, the advent of the Internet and the need to execute queries over streaming sources requires a different approach, since for streaming inputs the cardinality may not be known or may not even be knowable (as is the case for an unbounded stream.)

In view of this, we propose shifting from a cardinality-based approach ...

20 Research papers: adaptive, automatic, autonomic systems: Automatic physical database tuning: a relaxation-based approach



Nicolas Bruno, Surajit Chaudhuri

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Full text available: pdf(476.05 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

In recent years there has been considerable research on automated selection of physical design in database systems. In current solutions, candidate access paths are heuristically chosen based on the structure of each input query, and a subsequent bottom-up search is performed to identify the best overall configuration. To handle large workloads and multiple kinds of physical structures, recent techniques have become increasingly complex: they exhibit many special cases, shortcuts, and heuristics ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	2	"20040167891"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 08:56
L2	1	"6721724".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 08:59
L3	639	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (creat\$3 or generat\$3 or construct\$3) and group\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:06
L4	627	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:10
L5	24	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (directory near3 location\$2))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:08
L6	6	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (directory near3 location\$2) and (rank\$3 or unrank\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:08
L7	3	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (directory near3 location\$2) and (rank\$3 or unrank\$3) and (global same rank\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:15
L8	53	707/1-206.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (validat\$3 same plan\$2))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:12

L9	0	709/104.ccls. and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (validat\$3 same plan\$2))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:14
L10	36	Galindo-legaria\$.in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:14
L11	368	Waas\$.in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:15
L12	14	10 and 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:15
L13	3	12 and ((execut\$3 near3 plan\$2) and database\$2 and (search\$3 or quer\$3) and (creat\$3 or generat\$3 or construct\$3) and group\$2 and (directory near3 location\$2) and (rank\$3 or unrank\$3) and (global same rank\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/07/27 09:15